Machine Learning: Determining Fruit Ripeness from Visual and Auditory Data

(Technical Paper)

From Virtual Assistants to AI: Data Privacy Issues in the Digital Age

(STS Paper)

A Thesis Prospectus Submitted to the

Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements of the Degree

Bachelor of Science, School of Engineering

William Tan

Fall, 2022

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature _____ Date _____

William Tan

Approved _____ Date _____

Rosanne Vrugtman, Department of Computer Science

Approved _____ Date _____

Kathryn A. Neeley, Associate Professor of STS, Department of Engineering and Society

Prospectus Introduction

Today, virtual assistants are ubiquitous. Software like Siri, Alexa, and even Cortana can be commanded via speech, using a variety of speech recognition models in order to translate sound signals into words and then translate words into commands that can be followed. This software is often advertised as a provider of convenience since there is no other action required. There is no need to open a browser, open a new tab, and look through search results that may not even be relevant. However, all this convenience comes at a cost. In order to respond promptly, this kind of software must always be recording and capturing information that may be private. In addition, if data such as sound data is collected, what is inadvertently captured may also be of private or personal significance.

For the technical project, my objective is to determine if, despite noise in the foreground, other artifacts in the background, such as conversation, can be picked up and recognized. If background information can be captured, I will also see if there is a reliable way to obscure such information from various recognition algorithms. The goal is to see if there is a way to be able to only capture data that a user wants captured, and not take anything more than that.

As virtual assistants appear on more and more forms of technology, ostensibly for the sake of convenience, what is the true cost to the user? What privacy is the user giving up intentionally and unintentionally, and is the user aware of such a consequence? The information and decision-making that goes into the decision to use or not to use a virtual assistant, or some similar type of software can be complex, and I will explore the relationship between the user and the software, as well as the company that makes it. While users reduce their expectations of privacy when using this kind of software, they may be giving up more privacy than they expect

and it may fall upon the companies to mitigate this additional cost by informing users and modifying their software.

Technical Topic

As machine learning algorithms proliferate, all sorts of apps and software have begun to use it. This is especially true when it comes to processing audio and video information. By using the video feed from a camera, or the audio feed from a microphone, software that uses machine learning capture snapshots of the user's life and uses that information to fulfill its purpose. While it is mostly straightforward, being able to recognize songs, flowers, animals, and the like, it is what is captured inadvertently that can prove problematic. With the usage of software that relies on machine learning to process information becoming ubiquitous, we must assess the risk of recording potentially personal or sensitive information. In addition, users may agree for their data to be used for training purposes, and, while such data is ideally anonymized, it may still be possible to match a submission to its user. This prospect is possible if the pipeline for such data is compromised or if there is a backdoor.

Based in part on existing research in machine learning, my technical project will attempt to answer several questions based upon inadvertent information capture:

1. Is it possible to precisely and accurately separate foreground noise (which is intended to be captured) from background noise (which is often unintended to be captured)?
2. Is it possible to recognize data that is in the background?
3. How intact can background information be, and can anything be done about it?
4. Is it possible to reliably obscure background data to ensure that it cannot be captured intact before or after when it is recorded?

There exist models that already can take in a conversation and detect noises that occur in the background, able to attribute them to occurrences such as a dog barking. In addition, there also exist models that can reliably separate foreground noise, such as speech, from ambient noise. If the foreground noise can be isolated from the background, then isolating the background noise can also be attempted.

The environment and context in which data is recorded is also significant. While it may be impossible to determine where recording software is used, exactly, there probably exist surveys and other metrics that determine what setting it is used in, allowing for more insight into where and when most people use it. The level of awareness of a user should also be measured, since they could be completely oblivious to the fact that their device may be recording more things than they want it to. Lastly, since the controversy of voice assistants persists in users' minds, the attitude of the average user towards this potential concern would also be essential to determining an approach to educating and informing people.

Demonstrating the extent of information that can be inadvertently captured by software that uses machine learning for recognition is important to being able to reliably inform and educate users on the potential privacy concerns that may result from usage of this kind of software. I have already determined the aim of the technical project, which attempts to answer the questions listed above. Since there is already a surplus of models that aim to separate foreground noise from the data, it should be relatively intuitive to separate background noise from the data, as well. This should allow users to understand what they may be recording when using such software, and allow them to use it in a more educated manner, such as in a better environment.

STS Topic

Ever since the first voice assistants appeared on the market, there has been controversy to what it records (Solove 2006). Now that voice assistants can be voice-activated with a cuing phrase such as “Hey Siri!”, concerns have risen over how that phrase is detected. The phrase is detected because the voice assistant is constantly recording and transmitting, and what the software company does with the excess data has become a contentious topic. Allegations abound that the excess data is used for advertising and other algorithmic purposes, and the topic of privacy has loomed ever larger as voice assistants, as well as other types of software that almost constantly record, have become more and more commonplace. This makes it critical for the impact upon privacy, as well as any other such consequences, to be evaluated. Using the framework of social contracts, which are more intuitive to understand, I will analyze their assumed terms and the degree to which the parties involved comply with them.

A social contract, as posited by Locke and Hobbes, is essentially an agreement between the members of a society to endorse and comply with the society’s morals and rules, giving up some individual freedoms for the society’s protection of the others (D’Agostino 2021). The social contract, in this context, applies to both the users of the software and the companies that developed the software. The users agree to endorse and comply with the company’s terms in exchange for the usage of the software, surrendering some of their rights in exchange for convenience and better service. However, the basis on which the users make this agreement upon is worryingly not substantial. Users implicitly understand that their data is archived and categorized for further use, but can be unaware of the scope of such use and the scope of the data being collected, which often dwarfs their expectations (Waddell 2016).

While numerous anecdotes fill blogs and forums about users’ Alexa devices suggesting the ordering of products relevant to their conversations, investigations have demonstrated that

most such devices, such as iPhones, constantly record and transmit data to the companies that designed them. These kinds of findings influenced the adoption of data security and privacy regulations. The most famous examples in the last decade that were inspired by data privacy concerns are the European General Data Protection Regulation as well as the California Consumer Privacy Act, both of which allow users more leverage in determining how their data is collected and used (Lucarini 2020). Regulation such as this make the social contract between firms and users more transparent and easier to both understand and enforce. As such, my STS research question will be how the social contract between users of such software and the creators of such software has been affected by modern developments and regulation.

Privacy is a fundamental human want, and people reasonably expect some privacy even when agreeing to use software that may reduce it. However, when using this type of software, most users have absolutely no control of how it is handled. It is simply recorded and sent – only nominal measures are taken to ensure a modicum of privacy. The user has no say in where it's sent, how it's used, and who can see it. Ideally, it is lost in the tons and tons of data the company uses, but it still contains a wealth of potentially identifying information. This can include phone numbers, digital signatures and fingerprints, IP addresses, and other personal information. This harvested information goes to advertisers, analysts, and even influence campaigns like the infamous Cambridge Analytica (Fowler 2019).

Most companies that design this software cultivate a culture of apathy or necessity: the prevailing idea is that the loss of privacy is necessary in order to provide the service, and that it is worth the utter convenience it provides. However, the situation is never this black and white; companies can undertake a variety of measures to further protect data that is recorded. By considering the controversies around Apple and Amazon, the two companies that are arguably

most notorious when it comes to recording potentially personal and private data, we can find evidence of the current degradation of privacy. Whether its via exploitation of a loophole in software's terms of service, simple negligence, or willful ignorance, data that should be private is often still distributed to third parties for use (Fowler 2019).

However, the onus does not entirely have to be on the company; the user is not completely powerless. There do exist actions and practices that can mitigate the unwarranted reduction in privacy that can be entirely done on the user's end. For instance, one could wear makeup, wear neutral clothes, and practice other forms of techniques to minimize significant characteristics being recorded (Browne 2015). This, in addition to regulation, will be the subject of my research, which should attempt to find out how both government regulation as well as user action can affect both data security as well as encourage companies to adopt more privacy-respecting practices.

Next Steps

For the rest of the Fall semester, I will create and develop a prototype for the technical project. There exist many implementations of noise processing, and the goals I want my program to accomplish are already outlined in the questions above. I will also conduct surveys on those who have used software that follows the characteristics defined above, such as Alexa, Siri, and Shazam, and also attempt to aggregate the users' perceptions and practices to maintain privacy. While I am not going to a technical project for my capstone, I believe that by the Spring semester a sufficient foundation will exist to conduct further research.

References

Bolton T, Dargahi T, Belguith S, Al-Rakhami MS, Sodhro AH. On the Security and Privacy Challenges of Virtual Assistants. *Sensors (Basel)*. 2021 Mar 26;21(7):2312. doi: 10.3390/s21072312. PMID: 33810212; PMCID: PMC8036736.

Browne, S. (2015). *Dark matters: On the surveillance of blackness*. Duke University Press.

D'Agostino, Fred, Gerald Gaus, and John Thrasher, "Contemporary Approaches to the Social Contract", The Stanford Encyclopedia of Philosophy (Winter 2021 Edition), Edward N. Zalta (ed.)

Dhakal, Parashar & Damacharla, Praveen & Javaid, Ahmad & Devabhaktuni, Vijay. (2018). *Detection and Identification of Background Sounds to Improvise Voice Interface in Critical Environments*. 10.1109/ISSPIT.2018.8642755.

Fowler, G. A. (2019, August 16). *It's the middle of the night. do you know who your iPhone is talking to?* The Washington Post.

Lucarini, F. (2022, July 20). *GDPR vs CCPA: What are the main differences?* EUGDPRAcademy.

Solove, D. J. (2006). A Brief History of Information Privacy Law. *GW Law Faculty Publications & Other Works.*

Waddell, K. (2016, May 24). *Why Digital assistants are a privacy nightmare.* The Atlantic.